[0018] FIG. **1** illustrates an example authoring system **100** having features that illustrate examples aspects of the disclosure. The authoring system **100** includes a storage device **120** storing a master copy of a document **150**. In one embodiment, the storage device **120** may include, but not limited to, a server, a client computer, or other computing device. In another embodiment, the storage device **120** can include one or more storage devices (e.g., a network of computing devices).

[0019] The authoring system **100** also includes at least one client computing device **110** that is communicatively coupled to the storage device **120**. Each of the client computing devices **110** can edit the document **150** by creating a client copy **155** of the document **150** and editing the client copy **155**. The client copies **155** of the document **150** are synchronized when the client computing devices **110** periodically send to the storage device **120** updates to be shared with the other client computing devices and periodically obtain from the storage device **120** updates from other client computing devices.

[0020] As the term is used herein, a client computing device **110** includes any computing device that obtains a client copy of a document to be authored from a master copy of the document. The client computing device **110** can be different from the storage device **120** or can include a different client account implemented on the storage device **120**. In one embodiment, a computing device that acts as a storage device **120** for one document may act as a client computing device **110** for a different document and vice versa.

[0021] In the example shown, four client computing devices **110**A, **110**B, **110**C, and **110**D are communicatively coupled to the storage device **120**. In other embodiments, however, any number of computing devices **110** may be coupled to the storage device **120**. In the example shown, each client computing device **110**A, **110**B, **110**C, **110**D can send to the storage device **120** updates generated by the client of the client computing device and can request from the storage device **120** updates generated by the clients of the other client computing devices. In one embodiment, the storage device **120** can be a server computing device and the client computing devices **110**A, **110**B, **110**C, **110**D can be client computing devices. Other system configurations are possible. For example, in an alternative embodiment, multiple server computing devices can be used.

[0022] As shown in FIG. **2**, the document **150** stored on the storage device **120** can include content **152** and metadata **154**. Authoring applications **130** on the client computing devices **110** process and manipulate the content and metadata of the client copies **155** of the document **150**. In some embodiments, metadata **154** can be stored separately from content **152**. For example, content **152** can be stored in the document **150** and metadata **154** can be stored in a table (see FIG. **3**) separate from the document **150**. In other embodiments, however, the metadata **154** can be stored within the document **150**.

[0023] In general, the client computing devices **110** can synchronize updates to the content **152** separately from updates to the metadata **154**. In general, metadata updates **154** are automatically synchronized among the storage device **120** and client computing devices **110**, whereas content updates **152** from each client computing device **110** are synchronized at the request of the respective client.

[0024] Referring to FIG. **3**, lock metadata can be stored in a variety of different formats. For example, the lock metadata of FIG. **3** is stored in a table format **300**. The lock table **300** of FIG. **3** includes a list of clients, each of whom is identified with a client identifier (e.g., an identification number) that is uniquely assigned to the client. Data units to be locked are identified with unit identifiers (e.g., identification numbers) that are uniquely assigned to each data unit within a document. The lock table **300** associates the unit identifiers of the one or more data units to be locked with the client identifiers of the clients who own the locks.

[0025] For example, in the lock table **300**, data units **312** and **314** are associated with a first client **310**. Other clients, therefore, are inhibited from editing data units **312** and **314**. Data unit **322** is associated with client **320**. Other clients, including the first client **310**, therefore, are inhibited from editing data unit **322**. The fourth client **340** has not locked any portion of the document and so is not associated with any unit identifiers. In other embodiments, however, lock metadata can be stored in a different format or within the document. For example, the lock table **300** can be arranged by unit identifier instead of by client identifier.

[0026] Presence metadata also can be stored in a variety of formats. For example, presence metadata can be stored in the lock table **300** of FIG. **3**. In another embodiment, however, presence metadata can be stored in a separate table or in a different format. Presence metadata includes the client identifier of each client that is currently accessing the document or that has staked a claim (e.g., generated a content lock) on a data unit of the document. For example, a metadata table, such as the lock table **300**, can store the client identifier of each client having a claim to at least one data unit of the document. Like lock metadata, presence metadata can be synchronized automatically.

[0027] FIGS. **4** and **5** provide greater detail in how coauthoring between the client copy and the master copy of the document is implemented by a client computing device. FIG. **4** is a schematic block diagram of an authoring system **400** including a storage device **420** on which a master copy of a document to be authored is to be stored. The authoring system **400** also includes at least one client computing device **410** communicatively coupled to the storage device **420**.

[0028] The client computing device **410** includes an authoring application **412** configured to provide an authoring environment in which a client can create and/or manipulate a document to be authored. The client computing device **410** also includes a cache **414**, a layer object ("LO") **416**, and a synchronization manager ("sync manager") **418**. The cache **414** stores a client copy of the document to be authored. The cache **414** also stores the metadata, including lock and presence metadata, associated with the document. Updates to the content and metadata of the document also can be stored in the cache **414**.

[0029] The layer object **416** provides an interface between the authoring application **412** and the cache **414**. The layer object **416** also provides an interface between the authoring application **412** and the sync manager **418**. The sync manager **418** communicates with the storage device **420** and provides an interface between the storage device **420** and the cache **414**. For example, the sync manager **418** can send updates to and obtain updates from the storage device **420** and the cache **414**.

[0030] In general, an authoring environment having features that are examples of aspects in accordance with the principles of the disclosure can be implemented on a client computing device (e.g., a personal computer, a server computer, a notebook computer, a PDA, a Smartphone, or any